arXiv:1512.08133v1 [cs.LG] 26 Dec 2015

The Utility of Abstaining in Binary Classification

Akshay Balsubramani
University of California, San Diego
abalsubr@cs.ucsd.edu

Abstract

We explore the problem of binary classification in machine learning, with a twist - the classifier is allowed to abstain on any datum, professing ignorance about the true class label without committing to any prediction. This is directly motivated by applications like medical diagnosis and fraud risk assessment, in which incorrect predictions have potentially calamitous consequences. We focus on a recent spate of theoretically driven work in this area that characterizes how allowing abstentions can lead to fewer errors in very general settings. Two areas are highlighted: the surprising possibility of zero-error learning, and the fundamental tradeoff between predicting sufficiently often and avoiding incorrect predictions. We review efficient algorithms with provable guarantees for each of these areas. We also discuss connections to other scenarios, notably active learning, as they suggest promising directions of further inquiry in this emerging field.

1 Introduction

Consider a general practice physician treating a patient with unusual or ambiguous symptoms. The general practitioner often does not have the capability or equipment to confidently diagnose such an ailment, but will be able to refer the patient to a specialist or hospital. Therefore, the GP is faced with a difficult choice: either make a potentially erroneous diagnosis and act on it (which can occasionally have catastrophic consequences) or avoid committing to such a diagnosis and refer the patient on instead (which will almost certainly cost extra time and resources).

Such a situation motivates the study of prediction algorithms which are able not only to form a hypothesis about the correct prediction, but also *abstain* entirely from making a prediction if they are not confident enough. In machine learning, these algorithms are often formulated as classifiers, and the field of classification with abstention is therefore of increasing research interest.

In this manuscript, we survey the problem of binary classification where the classifier is allowed to abstain.

1.1 Outline

We will concern ourselves with the provable potential benefits of abstaining in two distinct scenarios. The first, covered in Section 2, is an online learning model free of distributional assumptions on the data, similar to the classical Mistake Bound model [Lit88]. The other, discussed in Section 3, is a standard statistical learning model in which the algorithm is initially given a chance to train on a set of labeled data, and the data are assumed to be drawn IID (independently and identically distributed) from a fixed distribution.

The abstain option in general, and each of the aforementioned models in particular, is intimately related to the well-studied paradigm of active learning, which aims to minimize the amount of labeled data required by a learner.

One idea that will recur in the survey, and is common to both settings we study, is the fundamental tradeoff between abstaining at a low rate and achieving low error when predictions are made. The intuitive reason for this is simple. A classifier will abstain exactly when it is unsure about which label to predict, and such unsureness implies that it would err on such examples relatively often if forced to predict on them. Traditionally, the classifier is indeed

forced to always predict; the possibility of doing better than this case is another spur to explore the utility of abstention.

Note that we avoid discussing Bayesian or explicitly distribution-dependent methods here [Cho70], as they are outside the scope of the manuscript; we believe the work surveyed here is itself extremely general and provides a unique web of insights into the problem.

Due to the volume of the material covered here and the necessary concision of this manuscript, full proofs of results will normally not be provided; the interested reader is referred to the appropriate references. However, we will endeavor to make the treatment here as standalone as possible, and provide proof sketches and intuitions whenever they are instructive.

1.2 Preliminaries

We consider a standard binary classification setting in which there exist unlabeled data points $x \in \mathcal{X}$, each of which is associated to a label in $\mathcal{Y} \triangleq \{-1, +1\}$; the pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ is referred to as a labeled data point. When the labels are deterministic, we can think of them as being generated by a map $h^* : \mathcal{X} \mapsto \mathcal{Y}$, so that for all labeled data points (x_i, y_i) we have $y_i = h^*(x_i)$.

The goal is to determine a classifier (a hypothesis) $h : \mathcal{X} \mapsto \{-1, +1, \perp\}$, where the classifier can either predict one of the two admissible labels or output $\perp$, which is referred to as *abstaining*. The latter can be interpreted as "don't know," a statement of ignorance about which label is correct without committing at all to either. In this case, the risk of predicting the label of the given datum incorrectly is deemed to be excessive. The understanding is that a lower penalty is paid with the $\perp$ prediction than an incorrect ± 1 prediction, which we call a *mistake*.

The learning algorithms we consider are given a hypothesis class $\mathcal{H}$, which is a set of hypotheses $h : \mathcal{X} \mapsto \mathcal{Y}$ that is used to constrain the set of possible solutions and for computational reasons. We will consider algorithms which choose a single $h \in \mathcal{H}$, as well as algorithms which choose a weighted majority vote of all $h \in \mathcal{H}$. We also follow standard terminology in calling problems with $h^* \in \mathcal{H}$ *realizable* or *separable*. The non-realizable scenario is often referred to as the *agnostic* case.

We defer more specific notation to the body of the manuscript, where it is presented as needed.

2 Abstaining in the KWIK Model

In this section, we describe a general model which has been influential in recent advances in the abstention literature. It will serve to provide useful insights into the possibilities afforded by abstention, even in a very general setting.

2.1 KWIK: Formulation

The KWIK ("Knows What It Knows") model [LLWS11] is a distribution-free way of formulating the online solvability of problems with *zero* incorrect predictions in total. Algorithms which accomplish this must be self-aware enough to recognize *all* points where they might make an incorrect prediction and output $\perp$ on them, and yet still predict (correctly) on a nontrivial number of examples. The name of the model refers to this self-awareness that it requires from algorithms, which distinguishes it from the otherwise similar Mistake Bound (MB) model.

We first review the salient features of the popular MB model as a useful reference point for further investigation.

2.1.1 The Mistake Bound Model

The Mistake Bound (MB) model [Lit88] for binary classification is a general framework for the study of online binary classification algorithms. In it, the learning algorithm receives unlabeled data points chosen one at a time in some arbitrary way (by an adversary in the worst case). At the beginning of round t, the algorithm has the hypothesis h_t, and sees an unlabeled point x_t. It then predicts a label $h_t(x_t) = \hat{y}_t \in \mathcal{Y}$, after which the true label y_t is revealed. The algorithm then makes any necessary internal modifications to update its hypothesis h_t to h_{t+1}. The next point x_{t+1} is then chosen, and the process repeats.

In this model, results of interest bound the total number of mistakes $|\{t : \hat{y}_t \neq y_t\}|$ made by the learning algorithm during the process, and are called *mistake bounds*. Typically, the true labels are assumed to be generated by some

hypothesis in $h^* \in \mathcal{H}$, so that $y_t = h^*(x_t)$. In this setting, learnability is defined as follows. Here $dim(\mathcal{H})$ represents some predefined notion of the complexity of $\mathcal{H}$ (for finite $\mathcal{H}$, take $dim(\mathcal{H}) = |\mathcal{H}|$).

Definition 2.1. *An algorithm learns $\mathcal{H}$ in the MB model if for any $\delta \in (0, 1)$, the algorithm makes a total of at most $\text{poly}(\frac{1}{\delta}, dim(\mathcal{H}))$ mistakes with probability at least $1 - \delta$ over entire runs of the algorithm.*

Well-known examples of mistake bounds include those for the perceptron [KWA97] and weighted majority [LW89] learning algorithms.

Such results are considered quite strong, since they imply that the algorithm eventually is mistake-free and because each unlabeled datum x_t may even be chosen by an adversary with full knowledge of the past $\{x_i, \hat{y}_i, y_i\}_{i=1}^{t-1}$. Guarantees in the standard setting of statistical learning, which we will consider more thoroughly in Section 3, can be seen to be less general primarily because they assume the unlabeled data x_t to be drawn IID from a distribution that is fixed a priori.

The KWIK model, however, is a strict restriction of this MB model, as we will see next.

2.1.2 The KWIK Model

The KWIK [LLWS11] model is used to analyze online supervised learning algorithms. The setting is very similar to the MB setting, with the algorithm seeing an adversarially chosen point x_t on the t^{th} round. However, the algorithm predicts $\hat{y}_t \in \mathcal{Y} \cup \{\perp\}$, and only observes the true label when it predicts $\perp$.

As in the previous discussion, we assume realizability: $h^* \in \mathcal{H}$. Then we can define our goal in this model, KWIK-learnability:

Definition 2.2. *An algorithm learns $\mathcal{H}$ in the KWIK model if for any $\epsilon, \delta \in (0, 1)$, the following are true with probability at least $1 - \delta$ over entire runs of the algorithm:*

1. *If $\hat{y}_t \neq \perp$, then $|\hat{y}_t - y_t| < \epsilon$.*
2. *$|\{t : \hat{y}_t = \perp\}|$, the total number of times the algorithm abstains, is at most $\text{poly}(\frac{1}{\epsilon}, \frac{1}{\delta}, dim(\mathcal{H}))$.*

It is clear that this shares many similarities with Definition 2.1. Some notes on Defn. 2.2 are in order:

- The first requirement corresponds to the algorithm making no mistakes when it does choose to predict, which is the hallmark of this framework. (Note: As pointed out by Li et al. [LLWS11], this does not make sense in the non-realizable case $h^* \notin \mathcal{H}$. Therefore, the definition can be generalized to the non-realizable case by relaxing the first requirement, though we only focus on the realizable case which most known results address.)
- The definition applies equally to regression problems (though we have presented y_t as being a discrete classification label). Another generalization is to the case when instead of observing the true label y_t, the algorithm observes a possibly randomized quantity z_t that depends on y_t. For instance, when learning the bias of a biased coin, z_t is a Bernoulli random variable with bias y_t. We neglect to formalize this for simplicity, but the full statement is in [LLWS11].

KWIK learnability is a fairly general condition for the same reasons discussed for MB learnability in Section 2.1.1. In fact, it is stronger:

Proposition 2.3. *Any KWIK algorithm which outputs $\perp$ at most B times can be converted into an MB algorithm for the same problem which makes at most B mistakes.*

Proof. The MB algorithm is the same as the KWIK algorithm, except that it outputs a prediction every time the KWIK algorithm would output $\perp$. □

As discussed earlier, this is due to the self-awareness required of the algorithm. An MB algorithm, by contrast, need only ensure it is learning correctly, as it blithely predicts without needing to consider its confidence in its own predictions.

For this reason, there are problems which are MB-learnable, but for which no KWIK algorithm exists; in a sense, the converse of Prop. 2.3 does not hold. One such problem is learning a Boolean singleton: the hypothesis class is $h_i : \{0, 1\}^n \mapsto \mathcal{Y},\ i = 1, \ldots, 2^n$, predicting $+1$ iff the input is a binary representation of $i - 1$. This can be

MB-learned with at most 1 mistake by simply predicting -1 always until the mistake is made. However, KWIK-learning it is more difficult - the restriction that the algorithm must be *absolutely sure* to predict is excessively strong, because in the worst case it is not sure until seeing $2^n - 1$ examples (effectively brute-force searching the space for the $+1$ data point). For the formal details, the reader is referred to Example 1 of [SZB10].

The generality of the KWIK framework allows it to be used even for algorithms with adaptive sampling strategies, in which there are complex dependencies between the data points sampled. The framework was originally motivated by applications in reinforcement learning, where such situations arise regularly. It has proved a useful tool in proving optimality guarantees for such algorithms; a further discussion is beyond the scope of this manuscript, but the interested reader is referred to [LLWS11, DLL09, WSDL09, SS11], and to [AADK13] for a recent application to the setting of multi-armed bandits.

2.2 KWIK: Algorithmic Building Blocks

As we have seen, the KWIK definition is extremely general. Remarkably, however, there are efficient KWIK algorithms for several nontrivial problems. These are for the most part derived from a few generic algorithms, which provide glimpses into the type of reasoning that is fruitful in this distribution-free setting. We therefore present some of these representative "building block" algorithms, along with their associated proof ideas.

2.2.1 Enumeration

The most ubiquitous and central such method, the *enumeration* algorithm, works in the realizable case when $\mathcal{H}$ is finite, and is laid out as Algorithm 1.

Algorithm 1 Enumeration

$\mathcal{V} \leftarrow \mathcal{H}$
for i = 1, 2, ... **do**
  if $h(x_i)$ is the same $\forall h \in \mathcal{V}$ **then**
    Choose any $h \in \mathcal{V}$ and predict $h(x_i)$
  else
    Output $\perp$ and receive label y_i
    $\mathcal{V} \leftarrow \{h \in \mathcal{V} : h(x_i) = y_i\}$
  if $|\mathcal{V}| = 1$ **then**
    Terminate with final output $\mathcal{V}$

The algorithm maintains a *version space* $\mathcal{V}$ of the hypotheses consistent with all data seen so far, paring down $\mathcal{V}$ appropriately as more data are revealed. This is a very natural strategy for making conservative predictions. The optimal hypothesis h^* is always in $\mathcal{V}$ by construction; consequently, any predicted label $\hat{y}_i$ always agrees with $h^*(x_i)$, which is the true label y_i by the realizability assumption.

Every time the algorithm abstains, at least one hypothesis in $\mathcal{V}$ is predicting the wrong label, and such hypotheses are removed each iteration. Therefore, each abstention leads to the removal of at least one hypothesis from $\mathcal{V}$, leading to a KWIK bound of $|\mathcal{H}| - 1$ abstentions.

There is a natural correspondence between predicting $\perp$ here and simply requesting a label - abstention may be thought of as a couched label request. Given this and the simplicity of the enumeration procedure, it is not surprising to find that Algorithm 1 appears independently in *active learning*, a supervised learning paradigm where labels are only provided upon the learner's request [Set12]. These and other active learning connections are discussed in Sec. 4.

Although the enumeration algorithm is appealingly general and simple, the major drawback is the same one confronted by the active learning literature: explicitly maintaining the version space by storing all its constituent hypotheses can be computationally intractable, requiring an exponential amount of memory!

Fortunately, as we will observe later in the manuscript, there are nontrivial special cases for which $\mathcal{V}$ can be implicitly maintained so that the operations of Algorithm 1 can be performed without actually storing $\mathcal{V}$, even when $\mathcal{H}$ is infinite.

2.2.2 Other Basic KWIK Algorithms

In order to highlight the utility of KWIK, we will briefly touch upon a few other basic algorithms in this framework. All of them are efficient - they run in polynomial time.

- *Learning the bias of a coin*: The algorithm simply abstains on the first T examples to get IID Bernoulli-distributed labels in order to estimate the bias, and thereafter predicts on every example according to the empirical estimate of the bias. Applying the Hoeffding bound yields a KWIK bound of $T \in \mathcal{O}\left(\epsilon^{-2}\ln(1/\delta)\right)$.
- *Learning the distribution of an n-sided die*: This is the multinomial-distribution analogue of the above coin-learning problem. It can be solved by reduction to n "one-vs.-rest" coin-learning problems, and the algorithm is the same as for the coin-learning case. A Chernoff bound analysis yields the KWIK bound $\mathcal{O}\left(n\epsilon^{-2}\ln(n/\delta)\right)$.
- *Linear regression with no noise in $\mathbb{R}^n$*: This is KWIK-learnable by a beautifully simple algorithm that exploits linear algebraic structure. The algorithm maintains a running training set of points x and their corresponding linear function values $f(x) = \theta \cdot x,\ \theta \in \mathbb{R}^n$ (unknown θ). For any new point x_t, the algorithm first checks if x_t is in the linear span of the points in the training set so far, which can be done efficiently by solving a system of linear equations. If so, $f(x_t)$ is an efficiently determinable linear combination of the function values in the training set, and the algorithm can predict with certainty. Otherwise, it abstains and adds $(x_t, f(x_t))$ to the training set. The training set cannot grow beyond size n, so the KWIK bound is n.
- *Linear regression with additive white noise in $\mathbb{R}^n$*: Though neither the proof nor the algorithm are as simple as for the zero-noise case, the algorithm predicts using the least-squares solution, or abstains by assessing a confidence measure based on the eigenvalue spectrum of the data. The KWIK bound is $\tilde{\mathcal{O}}\left(n^3\epsilon^{-4}\right)$. (Note: This setting has been studied independently for a much weaker *oblivious*, rather than adaptive, adversary; under this assumption, a much better bound of $\tilde{\mathcal{O}}\left(n\epsilon^{-2}\right)$ is proved in [CBGO09].)
- Algorithms also exist for certain generic combinations of KWIK-learnable scenarios. For instance, if $\mathcal{H}_i,\ i = 1, \ldots, n$ are KWIK-learnable, then so is $\cup_i \mathcal{H}_i$, even when observations are corrupted by white noise. Also KWIK-learnable are hypothesis classes whose domains and ranges are respectively Cartesian products of the domains and ranges of KWIK-learnable classes $\mathcal{H}_i$.

The intent here is to highlight that KWIK is a nontrivial model despite its extreme generality. We will now address a central issue of this survey - the tradeoff between errors and abstentions - by exploring the interesting consequences of allowing the learner to make some mistakes.

2.3 Allowing a Few Mistakes in KWIK

Though KWIK has proved to be quite useful in the past few years, its restriction that the learner make absolutely no mistakes is somewhat extreme in many circumstances. As we will see, allowing a few mistakes in the framework can drastically improve abstention rates.

2.3.1 Enumeration With a Few Mistakes

The cornerstone enumeration method of Algorithm 1 has recently been extended to allow up to k mistakes for some prespecified parameter k, again in an appealingly simple and generic manner [SZB10]. This newer algorithm, which we dub *relaxed enumeration* in our presentation, is outlined in Algorithm 2, and again is motivated by the realizable case for finite $\mathcal{H}$.

(Note: Here, the true label of a point is revealed regardless of whether the algorithm abstains or not. This is still consistent with the original KWIK framework of Algorithm 1 because of that algorithm's zero-mistakes property.)

Algorithm 2 has an instructive interpretation as a generalization of Algorithm 1. Like its predecessor, it tracks a version space $\mathcal{V}$ and terminates when enough data have arrived to pare this down to a singleton set containing h^*. When an unlabeled datum x_i arrives, a vote is taken over the hypotheses in $\mathcal{V}$. The margin of the vote is used as a measure of confidence in deciding whether or not to predict. The threshold s that is used to make this decision is decreased in such a way that at most k mistakes are made; this is done by keeping track of the number of mistakes m made so far.

Algorithm 2 Relaxed Enumeration

$\mathcal{V} \leftarrow \mathcal{H}, s \leftarrow |\mathcal{H}|^{\frac{k}{k+1}}, m \leftarrow 0$
for i = 1, 2, ... **do**
  $\mu \triangleq \min_{\lambda \in \{-1,+1\}} |\{h \in \mathcal{V} : h(x_i) = \lambda\}|$, so that μ is the number of hypotheses in $\mathcal{V}$ predicting the minority label
  if $\mu \leq s$ **then**
    Predict with the majority, i.e. predict $\hat{y}_i \triangleq \arg\max_{\lambda \in \{-1,+1\}} |\{h \in \mathcal{V} : h(x_i) = \lambda\}|$
  else
    Output $\perp$
  Receive label y_i
  $\mathcal{V} \leftarrow \{h \in \mathcal{V} : h(x_i) = y_i\}$
  if $\hat{y}_i \neq y_i$ and $\hat{y}_i \neq \perp$ **then**
    $m \leftarrow m + 1$
    $s \leftarrow |\mathcal{V}|^{\frac{k-m}{k+1-m}}$
  if $|\mathcal{V}| = 1$ **then**
    Terminate with final output $\mathcal{V}$

Due to the tolerance for a few mistakes, the KWIK bound on the number of $\perp$s is far better than that of Algorithm 1. The proof uses an elegant inductive argument which is worth noting in full:

Proposition 2.4 ([SZB10]). *Algorithm 2, in the realizable case, makes at most k mistakes and outputs $\perp$ at most $(k+1)\,|\mathcal{H}|^{\frac{1}{k+1}}$ times.*

Proof. The proof proceeds by induction in k. The first mistake made reduces $|\mathcal{V}|$ so that $|\mathcal{V}| < |\mathcal{H}|^{\frac{k}{k+1}}$. Since the effective size of the hypothesis class being considered at any point is $|\mathcal{V}|$, the inductive hypothesis is that after the first mistake, the algorithm will terminate after at most $k-1$ mistakes and $k\,|\mathcal{V}|^{\frac{1}{k}} < k\left(|\mathcal{H}|^{\frac{k}{k+1}}\right)^{\frac{1}{k}} = k\,|\mathcal{H}|^{\frac{1}{k+1}}$ abstentions.

The base case of the induction concerns the number of abstentions before the first mistake. In this regime, $s = |\mathcal{H}|^{\frac{k}{k+1}}$, and $\mu > s$, so there are at least $|\mathcal{H}|^{\frac{k}{k+1}}$ hypotheses removed from $\mathcal{V}$ every iteration for predicting the wrong label. Since $|\mathcal{V}| = |\mathcal{H}|$ at first, the algorithm abstains at most $\frac{|\mathcal{H}|}{|\mathcal{H}|^{\frac{k}{k+1}}} = |\mathcal{H}|^{\frac{1}{k+1}}$ times before the first mistake.

Combining this with the inductive hypothesis, the total number of $\perp$s is $\leq k\,|\mathcal{H}|^{\frac{1}{k+1}} + |\mathcal{H}|^{\frac{1}{k+1}} = (k+1)\,|\mathcal{H}|^{\frac{1}{k+1}}$, which completes the induction. □

In the original paper [SZB10], the analysis follows that of the "Egg Dropping Game," a common brainteaser, which can be found in [GF08]. It appears to this author that the analysis of [SZB10] can in fact be significantly tightened and enhanced using further results in [GF08], which utilize a more complex inductive argument. As this is unpublished, however, we limit ourselves to pointing it out as an open problem.

In any case, it is apparent that even for small k, allowing a few mistakes leads to vast improvements over Algorithm 1 in the dependence on $|\mathcal{H}|$ - from $\mathcal{O}(|\mathcal{H}|)$ to $\mathcal{O}\left(|\mathcal{H}|^{\frac{1}{k+1}}\right)$. For instance, allowing just $k = 1$ mistake yields a bound of $2\sqrt{|\mathcal{H}|}$.

This again highlights one of the most fundamental ideas in this survey - the tradeoff between predicting sufficiently often and avoiding prediction errors. Though zero-error learning is significant in its own right, it is a rather extreme condition that implies a similarly extreme cost of a mistake relative to the cost of abstaining. Most applications would be content to allow a few mistakes in return for a significantly more useful classifier (one that abstains less). This promise is exactly what Prop. 2.4 quantifies.

2.3.2 Tractable Relaxed Enumeration for Linear Separators

Algorithm 2 in general suffers from the same computational intractability as Algorithm 1, because it too appears to require the version space to be explicitly stored. However, this can be circumvented in useful special cases, including when learning monotone disjunctions and linear separators [SZB10]. We discuss the intuition for the latter, as it is instructive; for formal proofs, refer to [SZB10].

Concretely, let $\mathcal{X} = \mathbb{R}^d$ be the data space with the Euclidean norm, and let $\mathcal{H}$ be the set of linear separators $\{w \in \mathbb{R}^d : \|w\| = 1\}$, so that the prediction rule is $h_w(x) = \text{sgn}(w \cdot x)$. Denote the optimum separator by w^*. Assume the data are separable with margin $\gamma \triangleq \min_x \frac{w^* \cdot x}{\|x\|} > 0$ - this loosely corresponds to the realizability assumption of relaxed enumeration. This is the setup used for the well-known standard perceptron algorithm [KWA97]. If points $x_1, \ldots, x_n$ have already been seen, it will also be useful to think of the problem as a linear program with d variables (one for each coordinate of w) and n halfspace constraints $w \cdot x_i > 0$ for $\{i : y_i = +1\}$ and $w \cdot x_i < 0$ for $\{i : y_i = -1\}$.

[SZB10] show how to efficiently implement Algorithm 2 here using the following insight: the algorithm does not need to store $\mathcal{V}$, because its decision-making process only requires it to deal with $|\mathcal{V}|$ and estimate $|\{h \in \mathcal{V} : h(x_i) = \pm 1\}|$.

The first quantity, $|\mathcal{V}|$, is the volume of the feasible set of the linear program (LP) mentioned above, and it contracts significantly whenever a mistake is made. The γ-separability assumption implies that this feasible set cannot shrink below a certain volume, as there is some room around w^* where consistent separators can fall. Combining these two facts controls the number of mistakes made, and finishes the use of $|\mathcal{V}|$.

However, the algorithm still needs to estimate the quantities $|\{h \in \mathcal{V} : h(x_i) = \pm 1\}|$ for arbitrary x_i to decide whether or not to abstain. This is essentially the problem of estimating the relative volumes of two convex sets. Using Hit-and-Run [LV06], a tool for uniformly sampling from convex bodies, we can simply sample uniformly from the feasible set of the current LP, and see whether the sampled point is on the $+1$ or -1 side. Repeatedly sampling points essentially extends this idea to a Monte Carlo estimate of the relative volumes of $|\{h \in \mathcal{V} : h(x_i) = +1\}|$ and $|\{h \in \mathcal{V} : h(x_i) = -1\}|$, which the algorithm can use to make its decision.

It turns out that the number of points required can be managed to yield a polynomial-time algorithm which is equivalent to Algorithm 2 with high probability. Formalizing all this leads to the main result, which nicely complements the well-known perceptron mistake bound of $\mathcal{O}\left(\gamma^{-2}\right)$ [SS12]:

Proposition 2.5 ([SZB10]). *Assume data are linearly γ-separable in $\mathbb{R}^d$, and define $R = \left(1 + \frac{\sqrt{d}}{\gamma}\right)^d$. Then for any $k > 0$, a linear separator can be learned with at most k mistakes and $\mathcal{O}\left(R^{1/k} \log R\right)$ abstentions.*

Omitting logarithmic factors, this bound on $\perp$ is $\approx \mathcal{O}\left(\left(\frac{\sqrt{d}}{\gamma}\right)^{d/k}\right)$ for small γ. This implies the appealing fact that the mistake tolerance k needs only to scale linearly in the dimensionality d to achieve a fixed-exponent polynomial dependence on γ.

Note that although Algorithms 1 and 2 were formulated for finite $\mathcal{H}$ with the counting measure for simplicity, we have just sketched a tractable implementation for an infinite $\mathcal{H}$ with the Lebesgue measure. This type of efficient Monte Carlo sampling can widely broaden the potential usability of the algorithms; for instance, linear separators constitute a very useful $\mathcal{H}$ in applications.

3 Abstaining in Statistical Learning

We now change course somewhat to examine the effect of the abstain option in a standard *statistical learning* setting, where unlabeled data points x_i are drawn IID from a distribution that is fixed beforehand. This can be thought of as a variation of the PAC (Probably Approximately Correct) model which is standard in learning theory.

As mentioned earlier, this setting is more restrictive than the KWIK model. It imposes an extra IID assumption on the data generation process, and also treats the batch case in which we begin tracking the algorithm's performance after it has already seen m labels, rather than penalizing the algorithm for performance at the start of the learning process as well. As might be expected, these more restrictive assumptions lead to a wider range of stronger results than are known for the KWIK case. These results will now be introduced.

Formally, we now consider a standard batch learning setting where we have a training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$ of points drawn i.i.d. from a distribution $\mathcal{D}$. The learning algorithm trains on these points and thereafter sees no more labels. To emphasize the point, the learner will never even see the labels of points on which it predicts $\perp$ (recall that this was the case in KWIK).

Define the error of a hypothesis $\epsilon(h) = Pr_{\mathcal{D}}(h(x) \neq y)$; this is analogous to the number of mistakes in the KWIK framework. Similarly, rather than speaking of the total number of abstentions, the appropriate quantity to discuss will be the probability of abstentions over $\mathcal{D}$. We also will speak of the training error (error on S) $\hat{\epsilon}(h) = \frac{1}{m}\sum_{i=1}^{m} \mathbf{1}(h(x_i) \neq y_i)$.

We will otherwise follow the pattern laid out in our KWIK discussion of Section 2, presenting results first for the zero-error case and then addressing the error-abstention tradeoff.

3.1 An Algorithm for Zero-Error Learning

As we did in the KWIK overview of Section 2, we begin with the case when the algorithm is not allowed to make any erroneous predictions at all, and attempt to minimize the abstain probability. It is a nontrivial (and recent) result that this is even possible in a general setting. However, we will see how it can be done, and use the resulting algorithm as a point of departure to explore some powerful recent results in this vein.

Recall that this zero-error scenario's KWIK analogue (no mistakes) was addressed by Algorithm 1, the enumeration algorithm. It turns out that an almost identical algorithm can be used for the realizable case in the statistical learning setting as well. This was introduced in recent work as the Consistent Selective Strategy (CSS) [EYW10], described in Algorithm 3.

Algorithm 3 Consistent Selective Strategy (CSS)

Given: training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$
$\mathcal{V} \leftarrow \{h \in \mathcal{H} : h(x_j) = y_j \ \ \forall j = 1, \ldots, m\}$
for any unlabeled point x **do**
    if $h(x)$ is the same $\forall h \in \mathcal{V}$ **then**
        Choose any $h \in \mathcal{V}$ and predict $\hat{y} = h(x)$
    else
        Output $\perp$

CSS is essentially a batch version of Algorithm 1, the enumeration KWIK algorithm for the realizable case. It has zero error with certainty for the same reason as mentioned for Algorithm 1 - by construction, $h^* \in \mathcal{V}$, and so any prediction must agree with h^* and therefore be correct.

As with Algorithm 1, though, the significance of CSS lies in the fact that it does not abstain too often:

Proposition 3.1 (Thm. 8, [EYW10])**.** *For Algorithm 3,*

$$Pr_{\mathcal{D}}(\perp) \leq \frac{1}{m}\left((\ln 2)\min(|\mathcal{H}|, |\mathcal{X}|) + \ln\frac{1}{\delta}\right)$$

Proof Sketch. Two hypotheses are equivalent if they have the same predictions on all points in $\mathcal{X}$, so there are effectively $K = 2^{\min(|\mathcal{H}|, |\mathcal{X}|)}$ possible subsets of $\mathcal{H}$. Of these, the proof considers only the $k \leq K$ which are *not* unanimous on all the data (only on at most $1 - \epsilon$ of it, for some $\epsilon > 0$). For any such subset (call it G_i, $i \in \{1, \ldots, k\}$), it is very likely that not all m points in the training set are drawn from G_i. This implies that $\mathcal{V}$ is different from G_i.

In this way, it can be argued that $\mathcal{V}$ is different from all G_i. Since the G_i are characterized only by being unanimous on $\leq 1-\epsilon$ of the data, $\mathcal{V}$ must be unanimous on $\geq 1-\epsilon$ of the data. Since we predict exactly when $\mathcal{V}$ is unanimous, we must predict quite often, so we cannot abstain very often. □

The $\Theta(1/m)$ dependence of the abstain probability establishes CSS as a nontrivial algorithm. In fact, there is a significant sense in which it is optimal!

Proposition 3.2 (Thm. 7, [EYW10])**.** *In the realizable case, any classifier which has zero error with certainty on any $\mathcal{D}$ with any h^* satisfies the following condition: if it predicts (a non-$\perp$ label) on a point x, then so does CSS. Consequently, there is no such classifier with $Pr_{\mathcal{D}}(\perp)$ less than CSS.*

Proof Sketch. The argument is given for strategies which do not randomize the choice of whether to abstain on a given point. The proof is by contradiction, so assume there is a classifier A that has zero error for any data distribution, and a point $x_0 \notin S$ such that $A(x_0) \neq \perp$ and $CSS(x_0) = \perp$. By the construction of CSS, there is a hypothesis $h_2 \in \mathcal{V}$ such that $A(x_0) \neq h_2(x_0)$.

Having fixed these facts, we now formulate a new classification problem (using the same $\mathcal{H}$ and $\mathcal{X}$) to prove the contradiction. Consider a new distribution $\mathcal{D}_n$ over $\mathcal{X}$ which puts positive probability on each of the m points in S and on x_0, and zero probability on all other points in $\mathcal{X}$. Also define the new label-generating hypothesis $h_n^* \triangleq h_2$. Then with positive probability (say $\delta > 0$), an m-element training set drawn from $\mathcal{D}_n$ will be equal to S. Now A is deterministic and only depends on the training set, so A will still predict on x_0 with probability δ. But when it predicts on x_0, it will be erroneous because we established that $A(x_0) \neq h_2(x_0) = h_n^*(x_0)$.

So A does not have zero error with certainty, and we have our contradiction.

The argument for methods which randomize the abstain decision follows along the same lines. □

Of course, Prop. 3.2 does not preclude the existence of a classifier which abstains less than CSS but only has zero error with high probability. Nevertheless, the result establishes the conceptual importance of CSS as the prototypical zero-error learning algorithm in statistical learning. It is particularly notable that the argument carries over to randomized algorithms as well - all of them are trumped in this setting by the deterministic CSS.

This author is not aware of this optimality argument being made for the variety of other extant CSS-like algorithms, though the logic is potentially extensible to these algorithms. It could be a particularly useful perspective to mention in the KWIK case for the enumeration algorithm, for which a potential combinatorial extension of this result is currently being examined by this author as a byproduct of the investigations of this manuscript.

3.2 Zero-Error Learning in the Infinite-$\mathcal{H}$ Case

Interestingly, these favorable properties of CSS hold only for finite hypothesis classes; in fact, it is impossible in general for CSS to avoid abstaining everywhere when $\mathcal{H}$ is infinite, even in the realizable case. The counterexample construction used to show this is of independent interest, and will now be outlined.

The construction is schematically shown in Figure 3.1. $\mathcal{H}$ is the class of linear separators in $\mathbb{R}^2$, and the data are uniformly distributed on two nonintersecting arcs, with one containing only positive examples and the other only negative examples as shown in the figure. It is clear that this is a realizable problem, because any separator lying between the arcs that does not intersect either one will have zero error. However, for any finite training set the CSS version space $\mathcal{V}$ will only be unanimous for data in a convex polygon-like region inscribed in each arc; an example is the two shaded regions in Figure 3.1. Unfortunately, these regions only contain a finite number of data points which collectively have measure zero. So although CSS will predict correctly where it essays a prediction, it becomes a trivial algorithm which will almost surely abstain.

This construction is a valuable one because similar arguments to the above make it a "hard" situation for many algorithms that depend on version space constructions and/or unanimity between hypotheses, which arise in other fields like active learning [DHM07]. Loosely speaking, it has the property that $\mathcal{H}$ intuitively does not "fit" the data as represented; linear separators are not a natural choice for data whose intrinsic geometry is nonlinear. We will revisit this construction briefly later.

3.3 Characterizing the Abstain-Error Tradeoff

We now relax the requirement of zero error from the abstaining classifier and address how abstention rates can be reduced if a low error rate is allowed. The discussion here is the IID-setting analogue to that in Sec. 2.3.

The tradeoff is illustrated in Figure 3.2, in which it is termed the *risk-coverage* tradeoff. Coverage is simply the probability that an algorithm predicts rather than abstaining, and therefore is exactly the complement of the abstain probability $Pr_{\mathcal{D}}(\perp)$ we have used so far. Risk is a generalization of error probability to arbitrary loss functions. For our purposes hereafter, it can be taken as synonymous with error probability.

The tradeoff for any algorithm is schematically shown by the dotted curve in the figure: lower risk can only be achieved at the expense of coverage. The upper axis (with r^*) represents the case when the algorithm is not allowed to abstain; this is the scenario in standard binary classification, and r^* represents the risk of the algorithm in this case. The left axis represents the zero-error case that CSS addresses, and c^* is the coverage here. Most

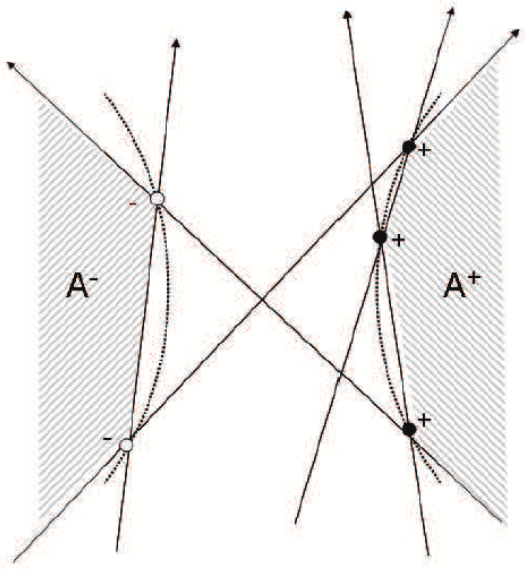

Figure 3.1: Construction for which nontrivial CSS zero-error learning is impossible (from [EYW10]).

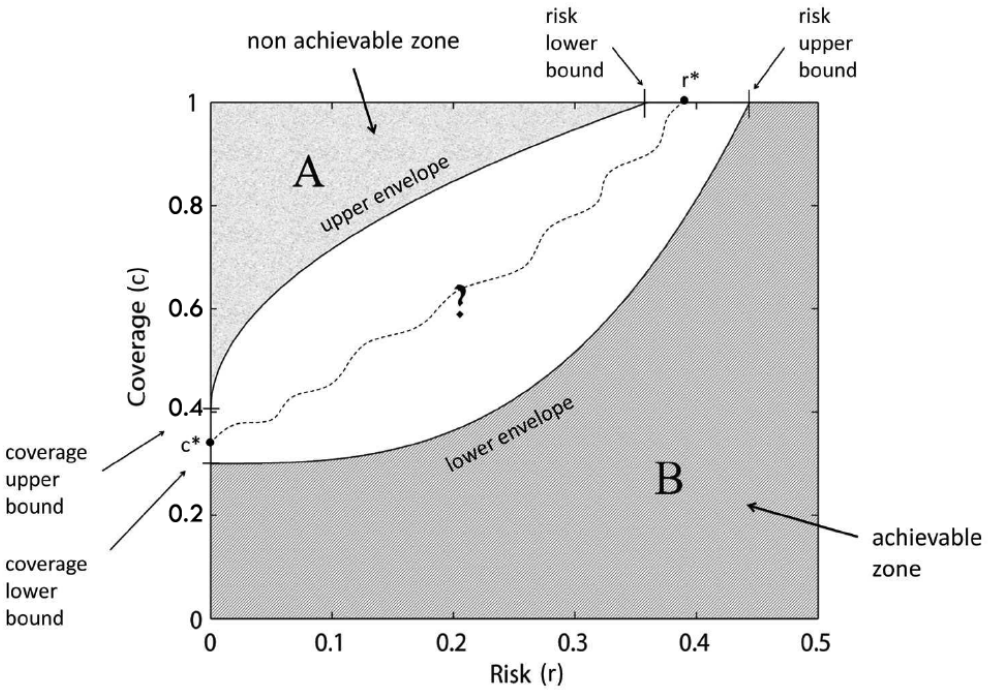

Figure 3.2: The risk-coverage tradeoff in statistical learning (from [EYW10]).

interestingly, upper and lower ”envelopes” on this tradeoff graph can be proved for abstaining classifiers, setting limits on the risk-coverage tradeoff in the statistical learning setting.

We present the upper envelope only, as it is the more fundamental result as proved in [EYW10].

Proposition 3.3 (Thm. 37 from [EYW10]). *Suppose d is the VC dimension of $\mathcal{H}$, and fix $\delta \in [0, \frac{1}{4}]$. Consider any classifier A with the abstain capability, and suppose it has abstain probability ρ and has error rate $R(A)$ on* examples on which it predicts. *Then there is a data distribution $\mathcal{D}$ such that with probability $\geq \delta$,*

$$R(A) \geq \min\left(\frac{1}{2} - \frac{1}{4\rho},\ \ \frac{1}{2} - \frac{1}{2\rho} + \frac{1}{16\rho m}\left(d + \frac{16}{3}\ln(1-2\delta)\right)\right)$$

No proof is given here for space reasons, but the reader is referred to [EYW10] to examine the interesting $\mathcal{D}$ construction and the associated VC dimension argument central to the proof.

The lower envelope proof involves interpolating between the two extremes of standard learning and CSS by considering a CSS-like classifier that predicts when $\mathcal{V}$ is unanimous, but flips a coin of bias $\alpha \in [0, 1]$ to decide whether to predict even when $\mathcal{V}$ is not unanimous. (Clearly, the $\alpha = 0$ case corresponds to CSS and $\alpha = 1$ to standard learning). The lower envelope is proved only for this α-parametrized classifier in [EYW10].

3.4 Relaxing CSS for the Non-Realizable Case

Though we have characterized the fundamental abstention tradeoff, our presentation so far has lacked an algorithm which, like Algorithm 2 in the KWIK setting, exploits the tradeoff. We use this section and Section 3.5 to explore two such methods in turn.

The first method we outline is directly inspired by CSS, and generalizes it to the non-realizable case. Here the target hypothesis $h^* = \arg\min_{h \in \mathcal{H}} \epsilon(h)$; the difference from the realizable case is that $\epsilon(h^*) \neq 0$ in general. The logic of CSS fails here, because h^* is not necessarily in the version space $\mathcal{V}$ as calculated on the training set S; indeed, it is possible that $\mathcal{V} = \emptyset$. However, the IID assumption on the data means that S is fairly representative of the data distribution $\mathcal{D}$. Therefore, if we relax the definition of $\mathcal{V}$ to include all low-training-error hypotheses (rather than zero-training-error hypotheses), we can trust that since $\hat{\epsilon}(h^*) \approx \epsilon(h^*)$, h^* will be in this newly defined $\mathcal{V}_{relaxed}$.

This is precisely what is done in [WEY11]. For convenience, we dub the resulting method "Relaxed CSS." It is given in Algorithm 4.

Algorithm 4 Relaxed CSS

Given: training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$, algorithm failure probability δ, VC dim. of $\mathcal{H} = d$
$\hat{h}^* \leftarrow \arg\min_{h \in \mathcal{H}} \hat{\epsilon}(h)$
$\mathcal{V}_{relaxed} \leftarrow \left\{h \in \mathcal{H} : \hat{\epsilon}(h) \leq \hat{\epsilon}\left(\hat{h}^*\right) + 4\sqrt{2\frac{d\ln(2me/d)+\ln(8/\delta)}{m}}\right\}$
for any unlabeled point x **do**
    if $h(x)$ is the same $\forall h \in \mathcal{V}_{relaxed}$ **then**
        Choose any $h \in \mathcal{V}_{relaxed}$ and predict $\hat{y} = h(x)$
    else
        Output $\perp$

This algorithm is shown in [WEY11] to satisfy good error and coverage bounds (with probability at least $1-\delta$ over the choice of S) if the "excess error" $\mathbf{1}(h(x) \neq y) - \mathbf{1}(h^*(x) \neq y)$ is well-behaved over $\mathcal{D}$ for all $h \in \mathcal{H}$. More precisely, its error on the data on which it predicts is at most the error of h^* on those data, with high probability. The full results are in [WEY11]; we omit them here due to space constraints in favor of discussing an interesting implementation trick which provides further connections to other work.

Algorithm 4, like CSS and both flavors of enumeration, faces the implementation difficulty of tracking a potentially exponentially large $\mathcal{V}_{relaxed}$ so that it can detect unanimity or lack thereof. However, this task can be reduced to a suitably chosen empirical error minimization, as indicated in [WEY11]:

Proposition 3.4 (Paraphrase of Lemma 6.1 in [WEY11]). *For any $x \in \mathcal{X}$, define $\hat{h}^*$ as in the algorithm and $\tilde{h}_x \triangleq \arg\min_{h \in \mathcal{H}}\{\hat{\epsilon}(h) \mid h(x) \neq \hat{h}^*(x)\}$. Then Algorithm 4 outputs $\perp$ iff $\hat{\epsilon}\left(\tilde{h}_x\right) - \hat{\epsilon}\left(\hat{h}^*\right) \leq 4\sqrt{2\frac{d\ln(2me/d)+\ln(8/\delta)}{m}}$.*

Proof. Working backwards, the condition $\hat{\epsilon}\left(\tilde{h}_x\right) - \hat{\epsilon}\left(\hat{h}^*\right) \leq 4\sqrt{2\frac{d\ln(2me/d)+\ln(8/\delta)}{m}}$ is equivalent to $\tilde{h}_x \in \mathcal{V}_{relaxed}$. By construction, $\tilde{h}_x(x) \neq \hat{h}^*(x)$. Both hypotheses are in $\mathcal{V}_{relaxed}$, so $\mathcal{V}_{relaxed}$ is not unanimous on x, and the algorithm therefore abstains. □

This simple argument reduces the generally intractable task of storing a version space to the task of empirical error minimization, which lacks the space complexity of the former and can be done efficiently in many cases (including for infinite $\mathcal{H}$). For this reason, it appears in several version space-based algorithms (including CSS, for which a variant is presented in [EYW10]). This is discussed in Section 4.2.

3.5 Aggregating Classifiers to Exploit the Tradeoff

We now turn our attention to a conceptually elegant algorithm that exploits the tradeoff in the statistical learning setting by taking a *weighted majority* (WM) vote of the classifiers in $\mathcal{H}$. We have so far explored two algorithms that generalize the zero-error strategy in different settings in an attempt to address the tradeoff: the relaxed CSS of Algorithm 4 and the relaxed enumeration strategy of Algorithm 2. Reexamining these will motivate the new WM voting method that we cover in this section.

Algorithm 4 generalizes CSS to the non-realizable case - essential to address the tradeoff - by relaxing the strict training set consistency requirement on the version space in CSS, instead considering the lowest-$\hat{\epsilon}(h)$ hypotheses in $\mathcal{H}$. Independently, Algorithm 2 generalizes enumeration to allow some mistakes in return for fewer abstentions - directly addressing the tradeoff - by relaxing the unanimity requirement on the version space for a prediction. Instead, it makes predict-or-abstain decisions based on the margin of a vote taken among hypotheses in the version space.

Combining these two separate relaxation ideas, a sensible algorithm could make predict-or-abstain decisions based on the margin of a vote over the lowest-$\hat{\epsilon}(h)$ hypotheses in $\mathcal{H}$, instead of assessing strict unanimity among only hypotheses with $\hat{\epsilon}(h) = 0$. This precise intuition can be realized by weighting each hypothesis h according to $\hat{\epsilon}(h)$ and then deciding whether to predict based on the WM vote. The resulting algorithm, described in [FMS04], is shown as Algorithm 5 for finite $\mathcal{H}$.

Algorithm 5 Weighted Majority Voting with Abstention

Given: training set $S = \{(x_1, y_1), \ldots, (x_m, y_m)\}$, failure probability δ, tuning parameter $\theta \in (0, 1/2)$
$\eta \leftarrow \ln(8|\mathcal{H}|)m^{\frac{1}{2}-\theta}$, $\Delta \leftarrow 2\sqrt{\frac{\ln(\sqrt{2}/\delta)}{m}} + \frac{\eta}{8m}$
for any unlabeled point x **do**
    $\hat{L}_\eta(x) \triangleq \frac{1}{\eta}\ln\left(\frac{\sum_{h:h(x)=+1} e^{-\eta\hat{\epsilon}(h)}}{\sum_{h:h(x)=-1} e^{-\eta\hat{\epsilon}(h)}}\right)$
    if $\left|\hat{L}_\eta(x)\right| > \Delta$ **then**
        Predict $\hat{y} = \mathrm{sgn}(\hat{L}_\eta(x))$
    else
        Output $\perp$

Here each hypothesis h are weighted as $e^{-\eta\hat{\epsilon}(h)}$, so that the high-weight hypotheses have low empirical error - an implicit encoding of the idea behind Algorithm 4, because only the high-weight hypotheses participate much in the vote. $\eta > 0$ is a learning rate that controls how much the weighting deviates from the default counting measure on $\mathcal{H}$; a larger training set leads to a larger η, implying greater confidence that the up-weighted hypotheses are indeed the ones with lowest true error $\epsilon(h)$. (To see this, consider the limit $m \to \infty$, in which case $\forall h \in \mathcal{H} : \hat{\epsilon}(h) \to \epsilon(h)$, and the weighting considers only $\arg\min_{h\in\mathcal{H}} \epsilon(h)$ while ignoring all other hypotheses). It can also be seen that the $\hat{L}_\eta(x)$ thresholding achieves the same effect as Algorithm 2: predict if the vote is enough of a landslide, and otherwise abstain.

As with the work we have previously examined, the theoretical guarantees proved for this algorithm include bounds on the abstain and error probabilities:

Proposition 3.5 (Corollary 1 in [FMS04]). *The error probability of Algorithm 5 is at most* $2\epsilon\left(h^*\right)+\tilde{\mathcal{O}}\left(m^{\theta-\frac{1}{2}}\right)+\delta$. *Also, for* $m \in \Omega\left(\left(\sqrt{\ln(1/\delta)}\ln(|\mathcal{H}|)\right)^{1/\theta}\right)$, *the abstain probability is at most* $5\epsilon\left(h^*\right)+\mathcal{O}\left(\frac{\sqrt{\ln(1/\delta)}+\ln(|\mathcal{H}|)}{m^{\frac{1}{2}-\theta}}\right)$.

Proof Sketch. The proof idea for the abstain bound is fairly standard and relies on concentration of $\hat{L}_\eta(x)$ about $L_\eta(x)$, its non-empirical analogue in which $\hat{\epsilon}\left(h\right)$ is replaced by $\epsilon\left(h\right)$ in the calculations. The abstain threshold Δ therefore can be seen to quantify the finite-sample uncertainty in estimating $L_\eta(x)$ by $\hat{L}_\eta(x)$.

The error bound, however, is proved differently. It essentially argues that the vote is dominated by "good" hypotheses ($\epsilon\left(h\right)$ close to $\epsilon\left(h^*\right)$), and uses the tautology that there must be at least one such good hypothesis - h^* itself - to argue that the vote too must perform well.

These arguments are quite general compared to other statistical learning arguments we have outlined, and hold equally well for the realizable and agnostic cases. They also hold for infinite $\mathcal{H}$ with an additional explicit assumption that there is a nonzero mass of "good" hypotheses. □

These bounds are seen to be quite loose (in their dependence on m) compared to those in Sec. 3.3. However, they do characterize the tradeoff, and exhibit a further interesting property: the error bound is completely independent of $\mathcal{H}$, as the dependence on hypothesis class complexity is corralled in the abstain probability. By comparison, the standard "Occam's Razor" IID-learning bound on a vanilla binary classifier $\epsilon\left(h^*\right)+\mathcal{O}\left(\sqrt{\frac{\ln(|\mathcal{H}|/\delta)}{m}}\right)$ does depend, albeit weakly, on $|\mathcal{H}|$ [BEHW87]. This is to our knowledge unique in the literature as another potential benefit of allowing the learner to abstain.

Though there is no explicit version space here, the predictions of all $h \in \mathcal{H}$ are nevertheless necessary, so the same intractability issues that plague CSS and enumeration recur here. This can be mitigated, for instance by carefully constructing a manageably finite "ϵ-cover" approximation of $\mathcal{H}$ [BB10]. However, it remains open how to use other tactics - such as the sampling and empirical error minimization ones previously mentioned - to address the problem.

The successive relaxations of unanimous-version-space arguments that we used to motivate this algorithm have the desirable consequence of making it much more robust. We illustrate this by noting that this WM algorithm is not foiled by the ill-matched geometry in the earlier construction of Figure 3.1. There are plenty of zero-error hypotheses in $\mathcal{H}$ in that scenario, and they dominate the vote, potentially leading to zero error for high enough m.

Votes like this are well known to be desirable in another respect: they can perform far better than h^* when the constituent hypotheses make diverse predictions, in effect compensating for each other's mistakes. This has further implications which are discussed in Sec. 5.

Finally, the WM algorithm used here has several other suggestive interpretations (for instance, if the total L^1 normalized weight of hypotheses predicting $+1$ on x is seen as an estimate of $Pr_{\mathcal{D}}\left(y=+1 \mid x\right)$, then $\hat{L}_\eta(x)$ is basically applying a logit transform). Other connections to previous work on similar weighted majorities in a variety of settings are too numerous to discuss here but highlight the richness of this approach [FS97, AHK12, LW94].

4 Connections to Active Learning

Active learning [Set12] is a well-motivated supervised learning paradigm in which the learner attempts to minimize the number of labels it needs by requesting labels at its own discretion. As mentioned in Section 2.2.1, there is a natural correspondence between a $\perp$ output and a label request, linking abstain-based results with active learning. We have previously interspersed brief references to these connections where appropriate, but will now devote this section to elaborating on how the work outlined in this manuscript relates to active learning.

A major goal of active learning in the realizable case is to learn a hypothesis of error ϵ using only $\mathcal{O}\left(\text{polylog}(1/\epsilon)\right)$ labels. This represents exponentially fewer labels than required by standard supervised learning, for which $\Omega(1/\epsilon)$ labels are needed [Das11]. This is an ambitious but sometimes achievable goal, as demonstrated by the canonical example of binary search as an active algorithm for learning a threshold on a line segment. It has also been proven for certain restricted special cases, like learning a linear separator when the data are uniformly distributed

on a sphere [FSST97, DKM09], but more general cases under which it is possible are much sought after in the community.

4.1 Zero-Error Learning

The online mistake-free enumeration algorithm (Alg. 1), developed for the KWIK setting, has been independently known as the CAL algorithm after the authors who first introduced it in [CAL94], and is the forerunner of a significant strand of theoretically motivated active learning work [BBL06, BDL09, BHLZ10]. This is somewhat surprising because it is extremely conservative in settling for $\perp$, whereas active learning aims to *minimize* the number of label requests. Indeed, one of the largest open problems in active learning is to design and analyze algorithms which are more "aggressive" in their label querying, rather than the more "mellow" CAL [Das11].

There is scant mention of this connection to active learning in the existing KWIK literature. However, we observed in Section 3 that the CSS algorithm is exactly analogous to enumeration in the IID-data setting of statistical learning. Recent work [EYW12] has shown a deep link between CSS and CAL (stated in terms of the VC dimension of $\mathcal{H}$, a widely-used measure of the hypothesis class complexity, although other such measures would also suffice):

Proposition 4.1 (Paraphrase of Thm.9 from [EYW12])**.** *Suppose $\mathcal{H}$ is a hypothesis class with VC dimension $d < \infty$. Then if CSS can learn it with $Pr_{\mathcal{D}}(\perp) \in \mathcal{O}\left(\frac{\text{polylog}(m)}{m}\right)$, CAL will learn it using only $\mathcal{O}(\text{polylog}(1/\epsilon))$ labels.*

This result motivates the further study of CSS to potentially yield highly desirable active learning results. A good beginning in this vein is also proved in [EYW12], using computational geometry tools from [EYW10] to show that CAL requests $\mathcal{O}(\text{polylog}(1/\epsilon))$ labels when $\mathcal{H}$ is the class of linear separators in $\mathbb{R}^d$ and $\mathcal{D}$ is a mixture of a fixed finite number of Gaussians. There remain concerns about how such approaches scale with the dimension d (an argument is made in [EYW12] that CAL must request exponentially many labels in d), indicating that CAL has its limitations. More aggressive query strategies may be needed.

4.2 Agnostic Learning

The idea of Algorithm 4 - to relax the strict training set consistency requirement of the version space in CSS and instead consider all hypothesis h with low enough $\hat{\epsilon}(h)$ - led to a major advance in active learning when it was successfully applied to CAL in [BBL06].

The implementation trick outlined earlier for Algorithm 4 - minimizing the empirical risk of a hypothetical dataset as a means of detecting unanimity among hypotheses in a relaxed version space - has also appeared several times in the abstention-related literature and in active learning [EYW10, SZB10, BHLZ10, DHM07, BBL06]. The "excess error" is assumed to be well-behaved in the analysis of the relaxed CSS algorithm; such assumptions on the same quantity have been partially studied in active learning as the Tsybakov noise conditions [Tsy04], under which efficient active learning algorithms with provably favorable properties have been derived. In fact, the excess error turns out to be fundamental to this line of active learning work, as made explicit recently in [ABE12].

5 Future Work and Conclusion

The study of abstaining classifiers is far from complete, and there are several open areas to be addressed. Prominent among them is the study of abstaining classifiers without a version space. The version space idea has been central to much work so far because of its thoroughly-explored utility in the realizable case when zero error is allowed. This is somewhat of a toy case compared to most real applications, however, and necessitates constant efforts to make the resulting algorithms more tractable. It is also an unnecessarily brittle notion, as the construction of Figure 3.1 suggests.

Another application-motivated open problem is to refine the insights of the work discussed here to be sensitive to the relative costs of abstaining and erring on a prediction. These costs generally drive the tradeoff in applications and could therefore be assumed to be known beforehand. An ambitious goal in this spirit is to devise algorithms which can precisely quantify their own confidence in a prediction, in a probabilistic manner.

Abstaining classifiers are also intimately connected to active learning, as discussed in Section 4. A range of ideas, from implementation tricks to algorithmic intuitions to the specific algorithms themselves, have been shown to recur in both fields. At present, there is very little work like [EYW12] that attempts to explicitly unify the fields

further. Both fields have a long history and have similar goals, yet there are several ideas, like relaxed enumeration in the KWIK setting and the disagreement coefficient in active learning [BBL06], that are currently explored asymmetrically in one field. Progress on cross-pollinating these ideas would be of mutual enrichment to both areas.

Another area of potentially fruitful research (on which we have made progress) concerns improvements on Algorithm 5, the weighted aggregation algorithm. We observed that this algorithm can be seen as synthesizing intuitions from relaxed enumeration and relaxed CSS, and so we believe the classifier aggregation approach for abstention is worth deeper examination. In particular, majority votes are remarkably robust, and are known to perform better than any of their constituents when the constituent predictions are diverse [KW03]. Hypothesis weighting strategies that maximize this diversity could provably help performance [WLR06], both for abstaining classifiers and in active learning.

We hope this survey provides some insight into recent approaches to designing classifiers that can abstain, which should see increased deployment in various applications over the coming years.

Acknowledgments

I would like to thank my committee - David Kriegman, Sanjoy Dasgupta, and Kamalika Chaudhuri - for their time and feedback at various stages of the research exam process. Thanks are also due to Yoav Freund for first introducing me to this area and for various helpful suggestions.

References

[AADK13] Jacob Abernethy, Kareem Amin, Moez Draief, and Michael Kearns. Large-scale bandit problems and kwik learning. In *ICML*, 2013.

[ABE12] Nir Ailon, Ron Begleiter, and Esther Ezra. Active learning using smooth relative regret approximations with applications. *Journal of Machine Learning Research - Proceedings Track*, 23:19.1–19.20, 2012.

[AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.

[BB10] Maria-Florina Balcan and Avrim Blum. A discriminative model for semi-supervised learning. *J. ACM*, 57(3), 2010.

[BBL06] Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Agnostic active learning. In Cohen and Moore [CM06], pages 65–72.

[BDL09] Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. Importance weighted active learning. In Danyluk et al. [DBL09], page 7.

[BEHW87] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Occam's razor. *Inf. Process. Lett.*, 24(6):377–380, 1987.

[BHLZ10] Alina Beygelzimer, Daniel Hsu, John Langford, and Tong Zhang. Agnostic active learning without constraints. In *NIPS*, pages 199–207, 2010.

[CAL94] David A. Cohn, Les E. Atlas, and Richard E. Ladner. Improving generalization with active learning. *Machine Learning*, 15(2):201–221, 1994.

[CBGO09] Nicolò Cesa-Bianchi, Claudio Gentile, and Francesco Orabona. Robust bounds for classification via selective sampling. In Danyluk et al. [DBL09], page 16.

[Cho70] C. Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory*, 16(1):41–46, 1970.

[CM06] William W. Cohen and Andrew Moore, editors. *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*. ACM, 2006.

[Das11] Sanjoy Dasgupta. Two faces of active learning. *Theor. Comput. Sci.*, 412(19):1767–1781, 2011.

[DBL09] Andrea Pohoreckyj Danyluk, Léon Bottou, and Michael L. Littman, editors. *Proceedings of the 26th Annual International Conference on Machine Learning, ICML 2009, Montreal, Quebec, Canada, June 14-18, 2009*, volume 382 of *ACM International Conference Proceeding Series*. ACM, 2009.

[DHM07] Sanjoy Dasgupta, Daniel Hsu, and Claire Monteleoni. A general agnostic active learning algorithm. In John C. Platt, Daphne Koller, Yoram Singer, and Sam T. Roweis, editors, *NIPS*. Curran Associates, Inc., 2007.

[DKM09] Sanjoy Dasgupta, Adam Tauman Kalai, and Claire Monteleoni. Analysis of perceptron-based active learning. *Journal of Machine Learning Research*, 10:281–299, 2009.

[DLL09] Carlos Diuk, Lihong Li, and Bethany R. Leffler. The adaptive k-meteorologists problem and its application to structure learning and feature selection in reinforcement learning. In Danyluk et al. [DBL09], page 32.

[EYW10] Ran El-Yaniv and Yair Wiener. On the foundations of noise-free selective classification. *Journal of Machine Learning Research*, 11:1605–1641, 2010.

[EYW12] Ran El-Yaniv and Yair Wiener. Active learning via perfect selective classification. *Journal of Machine Learning Research*, 13:255–279, 2012.

[FMS04] Yoav Freund, Yishay Mansour, and Robert Schapire. Generalization bounds for averaged classifiers. *The Annals of Statistics*, 32:1698–1722, 2004.

[FS97] Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, 1997.

[FSST97] Yoav Freund, H. Sebastian Seung, Eli Shamir, and Naftali Tishby. Selective sampling using the query by committee algorithm. *Machine Learning*, 28(2-3):133–168, 1997.

[GF08] William Gasarch and Stuart Fletcher. The egg game. *www.cs.umd.edu/ gasarch/BLOGPAPERS/egg.pdf*, 2008.

[KW03] Ludmila I. Kuncheva and Christopher J. Whitaker. Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy. *Machine Learning*, 51(2):181–207, 2003.

[KWA97] Jyrki Kivinen, Manfred K. Warmuth, and Peter Auer. The perceptron algorithm versus winnow: Linear versus logarithmic mistake bounds when few input variables are relevant (technical note). *Artif. Intell.*, 97(1-2):325–343, 1997.

[Lit88] Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. In *Machine Learning*, pages 285–318, 1988.

[LLWS11] Lihong Li, Michael L. Littman, Thomas J. Walsh, and Alexander L. Strehl. Knows what it knows: a framework for self-aware learning. *Machine Learning*, 82(3):399–443, 2011.

[LV06] László Lovász and Santosh Vempala. Hit-and-run from a corner. *SIAM J. Comput.*, 35(4):985–1005, 2006.

[LW89] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. In *FOCS*, pages 256–261. IEEE Computer Society, 1989.

[LW94] Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Inf. Comput.*, 108(2):212–261, 1994.

[Set12] Burr Settles. *Active Learning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool, 2012.

[SS11] István Szita and Csaba Szepesvári. Agnostic kwik learning and efficient approximate reinforcement learning. *Journal of Machine Learning Research - Proceedings Track*, 19:739–772, 2011.

[SS12] Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.

[SZB10] Amin Sayedi, Morteza Zadimoghaddam, and Avrim Blum. Trading off mistakes and don't-know predictions. In John D. Lafferty, Christopher K. I. Williams, John Shawe-Taylor, Richard S. Zemel, and Aron Culotta, editors, *NIPS*, pages 2092–2100. Curran Associates, Inc., 2010.

[Tsy04] Alexandre Tsybakov. Optimal aggregation of classiers in statistical learning. *The Annals of Statistics*, 32:135–166, 2004.

[WEY11] Yair Wiener and Ran El-Yaniv. Agnostic selective classification. In John Shawe-Taylor, Richard S. Zemel, Peter L. Bartlett, Fernando C. N. Pereira, and Kilian Q. Weinberger, editors, *NIPS*, pages 1665–1673, 2011.

[WLR06] Manfred K. Warmuth, Jun Liao, and Gunnar Rätsch. Totally corrective boosting algorithms that maximize the margin. In Cohen and Moore [CM06], pages 1001–1008.

[WSDL09] Thomas J. Walsh, Istvan Szita, Carlos Diuk, and Michael L. Littman. Exploring compact reinforcement-learning representations with linear regression. In Jeff Bilmes and Andrew Y. Ng, editors, *UAI*, pages 591–598. AUAI Press, 2009.